

# Revamping Multilingual Agreement Bidirectionally via Switched Back-translation for Multilingual Neural Machine Translation\*

Hongyuan Lu<sup>♡†</sup>, Haoyang Huang<sup>♣</sup>, Dongdong Zhang<sup>♣</sup>,  
Furu Wei<sup>♣</sup>, Wai Lam<sup>♡</sup>

<sup>♡</sup>The Chinese University of Hong Kong

<sup>♣</sup>Microsoft Corporation

{hylu,wlam}@se.cuhk.edu.hk

{haohua,dozhang,fuwei}@microsoft.com

## Abstract

Despite the fact that multilingual agreement (MA) has shown its importance for multilingual neural machine translation (MNMT), current methodologies in the field have two shortages: (i) require parallel data between multiple language pairs, which is not always realistic and (ii) optimize the agreement in an ambiguous direction, which hampers the translation performance. We present **Bidirectional Multilingual Agreement via Switched Back-translation (BMA-SBT)**, a novel and universal multilingual agreement framework for fine-tuning pre-trained MNMT models, which (i) exempts the need for aforementioned parallel data by using a novel method called switched BT that creates synthetic text written in another source language using the translation target and (ii) optimizes the agreement bidirectionally with the Kullback-Leibler Divergence loss. Experiments indicate that BMA-SBT clearly improves the strong baselines on the task of MNMT with three benchmarks: TED Talks, News, and Europarl. In-depth analyzes indicate that BMA-SBT brings additive improvements to the conventional BT method.

## 1 Introduction

Conventional multilingual neural machine translation (MNMT) leverages independent parallel data during the training process. In comparison, the multilingual agreement (MA) explicitly minimizes the output difference between two source inputs written in different languages but with the same meaning. Despite its success in from-scratch training on MT (Yang et al., 2021c), current methodologies suffer from at least two disadvantages that limit their scope of usage. Firstly, conventional

\*This research/paper was partially supported by the Center for Perceptual and Interactive Intelligence (CPII) Ltd. under the Innovation and Technology Commission’s InnoHK scheme.

<sup>†</sup>Contribution during an internship at Microsoft Research Asia.

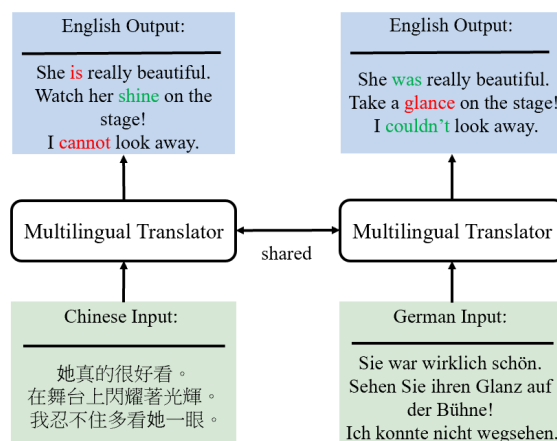


Figure 1: An illustrated example that can be benefited from Multilingual Agreement optimized in a bidirectional manner. The words in green are the correct translation, and the words in red are the wrong translation. Here, Chinese is incorrectly translated since it does not have past tense for verbs, and German is incorrectly translated due to the shared subword unit with different meanings between Glanz (German, shine) and Glance (English, take a brief look at). Best viewed in colour.

MA leverages word alignment tools to create code-switching sentence-level data (Yang et al., 2021c). This process usually requires authentic parallel data between multiple language pairs. For example, assuming we would like to enhance Chinese to English and German to English, conventional MA assumes the existence of parallel data from Chinese to German, which however sometimes does not exist. Secondly, the direction of agreement-based learning can be bidirectional (Zhang et al., 2019), while the direction of conventional multilingual agreement is usually ambiguous. However, since languages usually have different linguistic clues and they are helpful to each other, we argue that optimizing the multilingual agreement explicitly in a bidirectional manner can help the languages to learn from each other and hence further enhance cross-lingual learning.

Figure 1 depicts such a case that can be benefited from bidirectionally enhanced MA. The underlying reason is that both of the source inputs have cross-lingual ambiguities here. Since Chinese does not have past tense verbs, it is intuitive to use some auxiliary languages with past tense. Furthermore, since German shares partial vocabulary subwords with English under MNMT, this introduces cross-lingual ambiguities and using a language that does not share its vocabulary subwords with English, e.g., such as Chinese, could be helpful.

As a side note, since MA was proposed as a method for from-scratch training for MT, it was unclear whether conventional MA is also effective as a fine-tuning technique for pre-trained models.

Furthermore, how to appropriately apply back-translation to a multilingual setting is also an understudied subject despite its importance.

This paper proposes BMA-SBT, a novel MNMT framework that (i) exempts the need for parallel data between multiple language pairs and (ii) optimizes the MA in a bidirectional manner. To exempt the need for parallel data, we propose switched back-translation to produce synthetic text in some different auxiliary source languages with the translation target.<sup>1</sup> To optimize the MA in an explicit bidirectional manner, we use a bidirectional Kullback–Leibler Divergence loss instead of the code-switching for conventional MA. This enforces the original source language and the synthetic auxiliary language to have the same outputs as the target reference translation in a bidirectional manner.

We conduct experiments on three MT benchmarks: TED Talks (Cettolo et al., 2015), News benchmark (News-commentary) and Europarl (Koehn, 2005). Experimental results indicate that BMA-SBT clearly improves the strong pre-trained baselines on all three benchmarks. In-depth analyses indicate that BMA-SBT effectively mitigates cross-lingual ambiguities.

In summary, we make three key contributions:

- This paper proposes a novel framework called BMA-SBT, the first MNMT framework that achieves MA without the requirement of extra parallel data and explicitly optimizes the MA in a bidirectional manner.
- BMA-SBT yields clear improvement on SOTA pre-trained MT model on three MT benchmarks: TED Talks, News, and Europarl.

<sup>1</sup>For example, Chinese as the source, English as the target, and Japanese as the auxiliary source.

- We conduct in-depth analyses of BMA-SBT. Results indicate that BMA-SBT brings additive improvement to conventional BT and bidirectionality is important for MA.

Also, this is the first work that demonstrates the usefulness of MA as a fine-tuning technique.

## 2 Bidirectional Multilingual Agreement via Switched Back-translation

### 2.1 Multilingual Neural Machine Translation

We conduct our experiments on the task of MNMT on large-scale pre-trained multilingual translation model (Yang et al., 2021a; Lu et al., 2023) that handles multiple languages by sharing a universal subword dictionary among all the languages. For both training and inference, given  $I$  languages  $\{L_1, \dots, L_I\}$ , we prefix a special target language token  $L_t$  to the source inputs to signal the multilingual model that we are translating from an arbitrary source language to the target language  $L_t$ .

Given a bilingual dataset for machine translation that consists of  $\mathcal{N}$  training instances  $\{\mathcal{T}_1, \dots, \mathcal{T}_\mathcal{N}\}$ , each of the bilingual translation pairs  $\mathcal{T}_i$  in the source bilingual dataset  $\mathcal{D}_\mathcal{M}$  contains a source input  $x$  and the corresponding translation target  $y$ . With a Seq2Seq generation model (Sutskever et al., 2014) with parameters  $\theta$ , we train the model by optimizing the following likelihood:

$$\mathcal{L}_{main} = \sum_{n=1}^{\mathcal{N}} \mathbb{E}_{x_n, y_n \in \mathcal{D}_\mathcal{M}} [-\log P_\theta(y | x)], \quad (1)$$

where  $\mathcal{L}_{main}$  denotes the standard training loss that we adopt for MNMT.

### 2.2 BMA-SBT

In this subsection, we introduce our novel framework **Bidirectional Multilingual Agreement via Back-translation (BMA-SBT)**. Compared to the conventional multilingual agreement, BMA-SBT exempts the need for parallel data and specifies the direction of the multilingual agreement in a bidirectional manner. We first introduce how we use BT to create synthetic parallel data which are appropriate for the use of the multilingual agreement, and we then introduce how to leverage KL divergence loss to make the multilingual agreement bidirectional.

**Switched Back-translation** The conventional multilingual agreement (MA) requires authentic parallel data, which could be commonly unrealistic

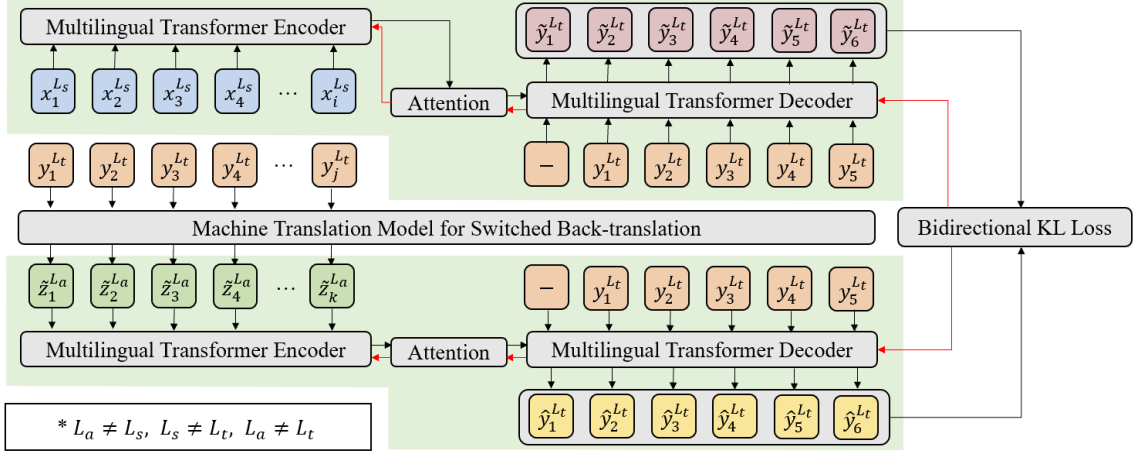


Figure 2: Overview of our proposed BMA-SBT framework.  $x$  and  $y$  denote the original source and target text written in the source language  $L_s$  and target language  $L_t$ .  $\tilde{z}$  denotes the synthetic text translated from the original target text into language  $L_a$ .  $\tilde{y}$  denotes the translation output from the original source text produced by the multilingual Transformer and  $\hat{y}$  denotes the translation output from the synthetic text. The letters with subscripts such as  $x_i$  denote the  $i$ -th token in the original source text. The red arrows denote the backward gradient flow computed by the bidirectional KL loss that updates the shared multilingual Transformer encoder and decoder. Best viewed in colour.

in a real-world setting. Formally, for the translation pair  $x$  and  $y$  in Equation 1, conventional MA requires another instance  $z$ , which is written in a different language and in the equivalent meaning to  $x$  and  $y$ . This process was designed and experimented on from-scratch training. These facts limit the use of the conventional MA.

To mitigate the above-mentioned shortages, we propose a novel method called switched back-translation that creates synthetic text  $\tilde{z}$  written in different source languages by feeding the translation target  $y$  to a machine translation model through back-translation.<sup>2</sup> Note that  $\tilde{z}$ ,  $x$ , and  $y$  are equivalent in their meanings, but they are written in different languages.

This helps us to establish a synthetic bilingual auxiliary dataset  $\mathcal{D}_A$  that is consisted of  $\mathcal{M}$  training instances. We then train the multilingual model by maximising the following likelihood:

$$\mathcal{L}_{auxiliary} = \sum_{n=1}^{\mathcal{M}} \mathbb{E}_{\tilde{z}_n, y_n \in \mathcal{D}_A} [-\log P_{\theta}(y | \tilde{z})]. \quad (2)$$

We also differentiate the switched back-translation we propose here from the conventional BT. For BT which was originally proposed for bilingual MT (Sennrich et al., 2016), we usually obtain

<sup>2</sup>While we can use the translation source  $x$  to create  $\tilde{z}$ , we empirically have found that this degrades the improvement. We postulate that if the source text has ambiguities, then this is less helpful to create the auxiliary text with the source text.

$x'$  from the original monolingual target  $y$ , where  $x'$  should be written in the same source language in our interest. In contrast, BMA-SBT creates  $\tilde{z}$  that should have the equivalent meaning as  $y$ , but it should be written in different languages from both the original source and target languages for the purpose of applying the multilingual agreement.<sup>3</sup>

In conclusion, this evolves the conventional MA into a universal fine-tuning technique for MNMT which does not need extra parallel data. BMA-SBT fits the real-world setting and can be applied with some modifications to other generation tasks for cross-lingual learning.

**Bidirectional Multilingual Agreement** The direction for agreement-based learning can be bidirectional (Zhang et al., 2019). However, the conventional multilingual agreement has an ambiguous direction due to the nature of code-switching. By using parallel data, conventional MA constructs code-switching data  $c$  from  $x$  and  $z$ , which denotes the translation source and the authentic auxiliary text respectively. Note that  $x$  and  $z$  have the same meaning to the translation target  $y$ , but they are written in different languages. The code-switching is then done with a word alignment tool between  $x$  and  $z$  at the word level, usually with a low code-switching replacement ratio as low as 10% (Yang

<sup>3</sup>For a fair comparison, we use the Baseline Model and the monolingual English sentences in the downstream dataset for data augmentation with BT (Sennrich et al., 2016) and SBT.

et al., 2021c). Formally, conventional MA trains MNMT by maximising the following likelihood:

$$\mathcal{L}_{MA} = \sum_{n=1}^Q \mathbb{E}_{c, y_n \in \mathcal{D}_C} [-\log P_\theta(y | c)], \quad (3)$$

where  $y$  denotes the translation target,  $\mathcal{D}_C$  denotes the code-switching dataset automatically constructed, and  $Q$  denotes the number of samples in the code-switching dataset.

In addition to the fact that conventional MA requires authentic data  $z$  which is not always realistic, we also argue that code-switching optimizes in an ambiguous direction, usually with a low code-switching ratio as low as 10%. Therefore, we consider that cross-lingual learning could be less efficient in this manner. As depicted in Figure 1, MNMT can be benefited by encouraging multilingual agreement in a bidirectional manner. Hence, we use a KL divergence loss to specify the direction of multilingual agreement in a clear bidirectional manner. Since the authentic parallel text  $z$  is not always available, we use the aforementioned synthetic auxiliary text  $\tilde{z}$  to calculate a bidirectional MA (BMA) divergence loss:

$$\mathcal{L}_{BMA} = \alpha \mathcal{L}_{KL_1} + (1 - \alpha) \mathcal{L}_{KL_2}, \quad (4)$$

where  $\mathcal{L}_{KL_1}$  and  $\mathcal{L}_{KL_2}$  represents the KL divergence loss in two directions:

$$\mathbb{E}[KL(P_\theta(y | x) || P_\theta(y | \tilde{z}))] \quad (5)$$

for  $\mathcal{L}_{KL_1}$ , which means that we enforce the original source text  $x$  to learn from the synthetic  $\tilde{z}$ . Note that  $x$  and  $\tilde{z}$  have the same meaning, but they are written in different languages. We also optimize in the other direction:

$$\mathbb{E}[KL(P_\theta(y | \tilde{z}) || P_\theta(y | x))] \quad (6)$$

for  $\mathcal{L}_{KL_2}$ .<sup>4</sup> In contrast to  $KL_1$ , this means that the synthetic text  $\tilde{z}$  should learn from the original text  $x$ . Bidirectionality is necessary to enforce both languages to learn from each other. Here,  $x$  and  $y$  denote the original translation source and target respectively, and  $\tilde{z}$  denotes the synthetic auxiliary text created by BMA-SBT via translation.

<sup>4</sup>Empirically, we have found that setting a balanced value with  $\alpha = 0.5$  brings a good performance.

**BMA-SBT** Overall, we propose a novel BMA-SBT framework that optimizes the MNMT models with the following combinatory loss:

$$\mathcal{L}_{BMA-SBT} = \mathcal{L}_{main} + \mathcal{L}_{auxiliary} + \mathcal{L}_{BMA} \quad (7)$$

Figure 2 depicts the overview of BMA-SBT. The final KL loss at the right edge of the figure refers to  $\mathcal{L}_{BMA}$ ,  $\mathcal{L}_{main}$  is calculated with the training instance at the top, and  $\mathcal{L}_{auxiliary}$  is calculated with the training instance at the bottom.

BMA-SBT can be improved with multiple auxiliary languages for agreement in an ensemble manner. This requires more tuning and computational costs. We leave this to future work.

## 3 Experiments

### 3.1 Implementation Details

**Model Configuration** The Transformer architecture we use is composed of 24 encoder layers and 12 interleaved decoder layers. Furthermore, the architecture has an embedding size of 1024, with a dropout rate of 0.1. The feed-forward network has a size of 4096, with 16 attention heads. For parameter initialization, we follow Ma et al. (2021) and Yang et al. (2021b) to pre-train a strong MT system with sentence-level bilingual data. For the rest of this paper, We call it the Baseline Model and use it as a strong pre-trained baseline system.

**Data Pre-processing** For all of the experiments conducted in this paper, we use SentencePiece (Kudo and Richardson, 2018) for tokenization. The SentencePiece model we use is the same as Yang et al. (2021b). Also, we follow prior works to prefix the source input translation texts with a language tag that indicates the target language of the outputs.

**Evaluations** We use the BLEU scores (Papineni et al., 2002) computed with the script from SacreBLEU for evaluation.<sup>5</sup>

**Training Details** We use the Adam optimizer (Kingma and Ba, 2014) and set it with the hyperparameter  $\beta_1 = 0.9$  and  $\beta_2 = 0.98$  for downstream fine-tuning. We set the learning rate as  $1e-5$ , with a warmup step of 4000. We use the label smoothing cross-entropy for the standard translation loss and we set label smoothing with a ratio of 0.1 for model training. All of the fine-tuning experiments reported in this paper are conducted on 8

<sup>5</sup><https://github.com/mjpost/sacrebleu>

Model	Fr→En	De→En	Zh→En	Vi→En	Cs→En	Th→En	Avg.
<i>Sentence-level Systems</i>							
HAN† (Miculicich et al., 2018)	-	-	24.00	-	-	-	-
M2M-100 (Fan et al., 2022)	50.18	42.24	26.62	34.92	37.84	27.28	36.51
mBART (Liu et al., 2020)	48.69	44.80	28.39	37.18	39.47	-	-
Baseline Model + BT	50.69	47.07	30.35	39.59	43.05	32.30	40.51
<i>Document-level Systems</i>							
mT5† (Xue et al., 2021a)	-	-	24.24	-	-	-	-
M2M-100 (Fan et al., 2022)	49.43	43.82	26.63	35.91	39.04	25.93	36.79
mBART (Liu et al., 2020)	49.16	44.86	29.60	37.09	39.64	-	-
MARGE† (Lewis et al., 2020)	-	-	28.40	-	-	-	-
Baseline Model + BT	49.53	45.98	30.17	39.28	42.33	30.62	39.65
Baseline Model + BT + MA (Yang et al., 2021c)	48.99	47.34	30.35	39.79	43.01	32.14	40.27
<i>Systems with Bilingual Parallel Document Data for Pre-training</i>							
DOCmT5†	-	-	31.40*	-	-	-	-
<b>BMA-SBT + BT</b>	<b>51.10</b>	<b>47.59</b>	<b>30.80</b>	<b>40.20</b>	<b>43.17</b>	<b>32.23</b>	<b>40.85</b>
<i>Ablation Study</i>							
- w/o $KL_1$	49.58	46.38	29.46	39.09	42.87	30.59	39.66
- w/o $KL_2$	50.56	47.47	30.26	40.02	43.15	31.89	40.56
- w/o $KL_1$ & $KL_2$	49.73	46.64	30.58	39.81	42.85	32.06	40.28

Table 1: Test results on TED Talks in the direction of (X → En). †: scores are taken from the official papers for these models. -: the scores are not reported or the language is not supported. \*: the score is not directly comparable due to the use of document-level parallel corpora for pre-training. The Baseline Model refers to the model described in Section 3.1, which is used for parameter initialization for BMA-SBT. BT refers to the conventional back-translation method described in Section 2.2.  $KL_1$  and  $KL_2$  refers to the loss described in Equation 5 and Equation 6 respectively. We train our system BMA-SBT at the document level.

NVIDIA V100 GPUs. We set the batch size as 512 tokens per GPU. Furthermore, to simulate a larger batch size, we update the models every 128 steps. For bilingual back-translation models, we use the downstream datasets for training on the same Transformer architecture.

### 3.2 TED Talks

**Experimental Settings** We use the IWSLT15 Campaign for the evaluation of TED Talks, on the task of multilingual MT. Prior systems have reported scores on only 1 or 2 translation directions (Lee et al., 2022; Sun et al., 2022), and Lee et al. (2022) supports only the translation direction into English (X → En). We report a wider range of language directions on the benchmark. We split all documents into sub-documents with a maximum of 512 tokens for all train/dev/test sets during training and inference. We use the official parallel training data from IWSLT15 with no additional monolingual data and the official 2010 dev set and 2010-2013 test set for evaluation (Liu et al., 2020; Lee et al., 2022). We use the Baseline Model to generate all the BT data and the SBT data used for multilingual agreement in BMA-SBT. We fine-tune

our model BMA-SBT at the document level. We report d-BLEU (Liu et al., 2020) using SacreBLEU.<sup>6</sup> d-BLEU score is a BLEU score for documents.

**Baseline Systems** We report strong baselines evaluated at both sentence and document levels. Evaluating at the sentence level means that we split documents into sentences for training and inference. In contrast, evaluating at the document level means that we split all documents into sub-documents with a maximum of 512 tokens as described in the Experimental Settings. We compare to the following baselines: M2M-100 (Fan et al., 2022), mBART (Liu et al., 2020), HAN† (Yang et al., 2016), MARGE† (Lewis et al., 2020), and the Baseline Model that we use to initialize the weights for BMA-SBT. †: the scores are taken from existing papers. We also report performance with Multilingual Agreement (Yang et al., 2021c) fine-tuned on Baseline Model with BT using synthetic parallel text. For a fair comparison, we do not directly compare to the SOTA model DOCmT5† (Lee et al., 2022), as it uses a large amount of bilingual parallel document data for a document-level multi-

<sup>6</sup><https://github.com/mjpost/sacrebleu>

Model	Fr→En	De→En	Zh→En	Cs →En	Avg.
<i>Sentence-level Systems</i>					
M2M-100 (Fan et al., 2022)	31.58	25.65	18.47	28.17	25.97
mBART (Liu et al., 2020)	29.93	29.31	18.33	30.15	26.93
<i>Document-level Systems</i>					
M2M-100 (Fan et al., 2022)	32.67	25.78	17.85	29.06	26.34
mBART (Liu et al., 2020)	30.14	26.35	15.01	29.79	25.32
Baseline Model (Yang et al., 2021b) + BT	36.38	34.24	25.58	36.97	33.29
<b>BMA-SBT (Ours) + BT</b>	<b>37.26</b>	<b>34.58</b>	<b>26.31</b>	<b>37.58</b>	<b>33.93</b>

Table 2: Test results on the News benchmark in the direction of ( $X \rightarrow \text{En}$ ).

lingual pre-training. The corpus used by DOCmT5 is not publicly available yet, and our methodology does not make use of such data. See Appendix A for the number of model parameters.

**Results** Table 1 presents the evaluation results of TED Talks in the direction of ( $X \rightarrow \text{En}$ ). BMA-SBT clearly surpasses the baselines. BMA-SBT surpasses the Baseline Model when both are fine-tuned at the document level by an average of 1.20 points in the score. BMA-SBT surpasses the Baseline Model fine-tuned at the sentence level by an average of 0.34 points in the score. Here, the Baseline Model fine-tuned at the document level is no better than that of the sentence level. We postulate that the underlying reason is that previous works have reported that directly optimizing the MNMT model at the document level can be challenging due to the long input problem (Koehn and Knowles, 2017). For a fair comparison, we add the conventional back-translation (BT) to both BMA-SBT and the Baseline Model. See Section 2.2 for more explanation on the difference between BT and the SBT methods used to achieve multilingual agreement.

In addition to the fact that BMA-SBT clearly improves the Baseline Model, which is a strong pre-trained MT system, BMA-SBT also beats other baselines such as HAN, M2M-100, mT5, and mBART, both fine-tuned at the sentence level and at the document level. Indeed, the Baseline Model itself is already quite competitive with these models, and being able to improve such a model is a piece of clear evidence for the effectiveness of BMA-SBT. The final results we obtain are close to the SOTA system DOCmT5, which uses a large amount of bilingual document translation pairs for multilingual pre-training.

**Ablation Study** The ablation study in Table 1 supports three points of view: (i) the bidirection-

ality of the multilingual agreement is necessary, (ii) the synthetic additional parallel data created by the BT used for MA is useful, and (iii) BMA-SBT brings additional improvements to the BT.

Firstly, the row of (-w/o  $KL_1$ ) and the row of (-w/o  $KL_2$ ) represent the ablations when the KL loss in the directions described in Equation 5 and Equation 6 are ablated respectively. Here, we can see that both lead to a degradation in the results. Clearly, using  $KL_2$  solely without  $KL_1$  seems to degrade the performance. This is not surprising, as  $KL_1$  pushes the output distributions of authentic data to be close to that of auxiliary text, which helps the model to use more linguistic clues in the auxiliary text. Also, using  $KL_2$  solely pushes the outputs of synthetic auxiliary data to be close to that of the authentic data unidirectionally, which can be less helpful to the original authentic data. Removing  $KL_2$  and using  $KL_1$  solely also degrades the results, which aligns with our original motivation depicted for the bidirectionality as in Figure 1.

Secondly, the row of (- w/o  $KL_1 \& KL_2$ ) brings improvements compared to Baseline Model + BT, which means that the auxiliary parallel data itself created by switched back-translation is useful.

Finally, BMA-SBT + BT brings clear improvements to the Baseline Model + BT. Since both models have used the conventional BT (See Section 3.1 for more details), the comparison is fair, which means that the BMA-SBT framework is effective and brings additive improvement to BT.

### 3.3 News

**Experimental Settings** For evaluation on the News benchmark, we use News Commentary v11 as the training set, following Sun et al. (2022). We employ newstest2015 as the dev set, and newstest2016/newstest2019 as the test set respectively for Cs and De. We use newstest2013 as the dev

Model	Da→En	De→En	El→En	Es→En	Fr→En	It→En	Nl→En	Pt→En	Sv→En
<i>Sentence-level Systems</i>									
M2M-100 (Fan et al., 2022)	50.40	47.38	52.28	52.03	48.26	49.70	46.78	49.84	52.34
Baseline Model + BT	48.94	47.25	53.46	50.57	47.68	49.49	45.95	50.65	52.77
<i>Document-level Systems</i>									
M2M-100 (Fan et al., 2022)	50.33	47.00	52.24	52.14	48.13	49.71	46.65	40.68	52.28
Baseline Model + BT	49.85	47.64	53.34	51.32	48.46	50.26	47.12	50.13	52.42
<b>BMA-SBT (Ours) + BT</b>	<b>50.52</b>	<b>47.86</b>	<b>54.06</b>	<b>52.17</b>	<b>48.77</b>	<b>50.67</b>	<b>47.90</b>	<b>50.69</b>	<b>52.96</b>

Table 3: Test results on the Europarl benchmark in the direction of (X → En).

Source	....., 当光在西红柿上走过时, 它一直在闪耀。它并没有变暗。为什么? 因为西红柿熟了, 并且光在西红柿内部反射, .....
Reference	..., as the light washes over <b>the tomato</b> , It continues to glow. It doesn't become dark. Why is that? Because <b>the tomato is actually ripe</b> , and <b>the light</b> is bouncing around inside the tomato, ...
Google Translate	..., as the light passed over <b>the tomatoes</b> , It kept shining. It didn't get darker. Why? Because <b>the tomatoes are ripe</b> , and <b>light</b> is reflected inside the tomatoes, ...
Microsoft Translator	..., as the light walks over <b>the tomatoes</b> , It keeps shining. It didn't darken. Why? Because <b>the tomatoes are ripe</b> , and <b>light</b> is reflected inside the tomatoes, ...
DeepL Translate	..., as the light traveled over <b>the tomatoes</b> , it kept shining. It doesn't dim. Why? Because <b>the tomatoes are ripe</b> and <b>the light</b> is reflecting inside the tomatoes, ...
Baseline Model (Sentence-level)	..., as the light goes over <b>the tomato</b> , It's always glowing. It's not darkening. Why? Because <b>the tomato is ripe</b> , and <b>light</b> is reflected inside the tomato, ...
Baseline Model (Document-level)	..., as the light passes over <b>the tomato</b> , It keeps flashing. It doesn't get darker. Why? Because <b>the tomatoes are ripe</b> , and <b>the light</b> is reflected inside the tomato, ...
BMA-SBT	..., as the light passes over <b>the tomato</b> , It's always shining. It's not darkening. Why? Because <b>the tomato is ripe</b> , and <b>the light</b> is reflected inside the tomato, ...

Table 4: A Chinese-to-English case study from TED Talks demonstrates that BMA-SBT captures better noun-related issues. We highlight the correct translation in cyan (the darker one when printed in B&W), and the mistakes in lime (the lighter one when printed in B&W). Google Translate: <https://translate.google.com/>, Microsoft Translator: <https://www.bing.com/translator>, DeepL Translate: <https://www.deepl.com/translator>. Time-stamped on 15th June 2023, can be subject to change.

set and newstest2015 as the test set for Fr. We use newstest2019 as the dev set and newstest2020 as the test set for Zh. The remaining settings follow the same as the evaluation on TED Talks.

**Baseline Systems** As the weights for DOCmT5 are not available at the time of writing, we compare our system to various strong baselines such as M2M-100, mBART and the Baseline Model. We run the fine-tuning process on the official checkpoints to obtain the scores. For a fair comparison, we apply BT to the Baseline Model.

**Results** Table 2 compares BMA-SBT to strong baselines, and we see that the improvements with BMA-SBT are clear, and the final results surpass all the strong baselines. This validates BMA-SBT's effectiveness as a novel framework.

### 3.4 Europarl

**Experimental Settings** For the Europarl dataset (Koehn, 2005), we use Europarl-v7 Sun et al. (2022). We experiment with (X → En) where we test nine languages: Da, De, El, Es, Fr, It, Nl, Pt, and Sv. Like previous works (Bao et al., 2021; Sun et al., 2022), the dataset is randomly partitioned

into train/dev/test divisions, and we split by English document IDs to avoid information leakage to better support the multilingual setting.

**Baseline Systems** As the weights for DOCmT5 are not available at the time of writing, we compare our system to various strong baselines such as M2M-100 and the Baseline Model. We run the fine-tuning process on the official checkpoints to obtain the scores. For a fair comparison, we apply BT to the Baseline Model.

**Results** Table 3 compares BMA-SBT to strong baselines, and we see that the improvements with BMA-SBT are obvious, and the final results surpass all the strong baselines.

### 3.5 Case Study

Table 4 depicts a Zh→En case study on TED Talks. In addition to the Baseline Models, we also compare BMA-SBT to various commercial systems such as Google Translate. In this case, we see that the Chinese text does not differentiate plural from single. Among all cases, it is clear that BMA-SBT works the best and can effectively resolve such ambiguity. We also observe that BMA-SBT perfectly

capture the context and attaches the definite article ‘the’ to ‘light’. This aligns with our original intention depicted in Figure 1 to help the models to improve cross-lingual learning via BMA-SBT.

### 3.6 Coherence and Consistency Evaluation

Figure 3 depicts the evaluations in the averaged scores from six translation directions on TED Talks with BlonDe scores (Jiang et al., 2022). BlonDe is an evaluation metric designed for MT which considers document-level coherence and consistency issues that require the model to resolve cross-lingual ambiguities. We see that BMA-SBT brings effective improvements to the metric.

## 4 Related Work

### 4.1 Multilingual Neural Machine Translation

Conventional bilingual machine translation models deal with two languages: one as the input, and one as the output. In comparison, multilingual neural machine translation (MNMT) has achieved great success in handling multiple languages with a single model. Recently, there have been many pre-training works on MNMT through multilingual pre-training models that leverage unsupervised pre-training objectives on monolingual corpora in many different languages (Conneau et al., 2020; Liu et al., 2020; Xue et al., 2021b). Following the calls that the unsupervised scenario is not strictly realistic for cross-lingual learning (Artetxe et al., 2020), subsequent works use parallel corpora with translation pairs for multilingual pre-training (Reid and Artetxe, 2022; Lee et al., 2022).

While pre-training has shown great success for MNMT (NLLB-Team, 2022), it is unclear whether the previous methods for from-scratching training on MNMT are still useful on pre-trained models. Multilingual agreement (Yang et al., 2021c) is perhaps the closest work to ours among those methods for from-scratching training. However, conventional MA requires authentic parallel data among many language pairs, which does not always guarantee to exist. In comparison, we focus on a more recent fine-tuning setting on popular pre-training models as well as a realistic setting with no presumption on the existence of the additional parallel data.

### 4.2 Agreement-based Learning

Agreement-based learning has been proven as a useful paradigm in the language community (Liang et al., 2006, 2007; Cheng et al., 2016). The core




Figure 3: Averaged BlonDe scores from six directions in ( $X \rightarrow \text{En}$ ) on the dataset of TED Talks evaluated with BMA-SBT and the Baseline Model (Document-level).

idea is to minimize the difference in the representations between the inputs with the same meaning. Some multilingual pre-training methods such as Chi et al. (2021) are relevant to agreement-based learning in the way that they shrink the distance of cross-lingual representations between parallel data. Zhang et al. (2019) proposed to enforce an agreement on the output with left-to-right and right-to-left inputs on recurrent neural networks for machine translation. Yang et al. (2020) proposed to use phrase-level agreement for machine translation.

Still, Yang et al. (2021c) is the closest work to ours, which encourages agreement between parallel data in different languages to have the same translation outputs. A very recent concurrent work uses MA to close the gap between source and target languages (Gao et al., 2023). Our work creates synthetic data and employs bidirectional KL loss to enforce the multilingual agreement bidirectionally.

## 5 Conclusions

Despite the fact that multilingual agreement (MA) has shown its effectiveness in from-scratching training for MNMT, the conventional MA has at least two shortages that limit its usages: (i) needs authentic extra parallel data, which can be often unrealistic and (ii) has an ambiguous direction for agreement-based learning. We propose BMA-SBT as a novel and universal fine-tuning framework for pre-trained MT models that (i) exempts the need for authentic parallel data by creating synthetic parallel text written in a different source language and (ii) specifies the direction of agreement-based learning with bidirectional KL divergence loss. Experimental results on three multilingual machine translation datasets illustrate that BMA-SBT can obviously improve the strong pre-trained baseline system. An in-depth investigation indicates that BMA-SBT brings additive improvements to the conventional BT methods for neural machine translation.



## Limitations

The proposed method requires generating synthetic auxiliary parallel data using translation models, which requires extra computational costs. The proposed method requires generating synthetic auxiliary parallel data using translation models, which requires extra computational costs.

**Large Language Models** Large language models (LLMs) such as ChatGPT have shown good translation abilities (Lu et al., 2023), while they still lag behind supervised systems (Jiao et al., 2023; Zhu et al., 2023). We do not directly compare them, as they are much larger in their number of parameters than the systems described in this work.

## Ethics Statement

We honour and support the EAACL Code of Ethics. The datasets used in this work are well-known and widely used, and the dataset pre-processing does not make use of any external textual resource. In our view, there is no known ethical issue. End-to-end pre-trained generators are also used, which are subjected to generating offensive context. But the above-mentioned issues are widely known to commonly exist for these models. Any content generated does not reflect the view of the authors.

## References

- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020. [A call for more rigor in unsupervised cross-lingual learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7375–7388, Online. Association for Computational Linguistics.
- Guangsheng Bao, Yue Zhang, Zhiyang Teng, Boxing Chen, and Weihua Luo. 2021. [G-transformer for document-level machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3442–3455, Online. Association for Computational Linguistics.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. [The IWSLT 2015 evaluation campaign](#). In *Proceedings of the 12th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 2–14, Da Nang, Vietnam.
- Yong Cheng, Shiqi Shen, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. Agreement-based joint training for bidirectional attention-based neural machine translation. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 2761–2767. AAAI Press.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. [InfoXLM: An information-theoretic framework for cross-lingual language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2022. Beyond english-centric multilingual machine translation. *J. Mach. Learn. Res.*, 22(1).
- Pengzhi Gao, Liwen Zhang, Zhongjun He, Hua Wu, and Haifeng Wang. 2023. [Improving Zero-shot Multilingual Neural Machine Translation by Leveraging Cross-lingual Consistency Regularization](#). *arXiv e-prints*, page arXiv:2305.07310.
- Yuchen Jiang, Tianyu Liu, Shuming Ma, Dongdong Zhang, Jian Yang, Haoyang Huang, Rico Sennrich, Ryan Cotterell, Mrinmaya Sachan, and Ming Zhou. 2022. [BlonDe: An automatic evaluation metric for document-level machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1550–1565, Seattle, United States. Association for Computational Linguistics.
- Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. [Is ChatGPT A Good Translator? Yes With GPT-4 As The Engine](#). *arXiv e-prints*, page arXiv:2301.08745.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.

- Philipp Koehn and Rebecca Knowles. 2017. **Six challenges for neural machine translation**. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. **SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Chia-Hsuan Lee, Aditya Siddhant, Viresh Ratnakar, and Melvin Johnson. 2022. **DOCmT5: Document-level pretraining of multilingual language models**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 425–437, Seattle, United States. Association for Computational Linguistics.
- Mike Lewis, Marjan Ghazvininejad, Gargi Ghosh, Armen Aghajanyan, Sida Wang, and Luke Zettlemoyer. 2020. Pre-training via paraphrasing. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA. Curran Associates Inc.
- Percy Liang, Dan Klein, and Michael I. Jordan. 2007. **Agreement-based learning**. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, pages 913–920. Curran Associates, Inc.
- Percy Liang, Ben Taskar, and Dan Klein. 2006. **Alignment by agreement**. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 104–111, New York City, USA. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. **Multilingual denoising pre-training for neural machine translation**. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Hongyuan Lu, Haoyang Huang, Shuming Ma, Dongdong Zhang, Wai Lam, Zhaochuan Gao, Anthony Aue, Arul Menezes, and Furu Wei. 2023. **TRIP: Accelerating document-level multilingual pre-training via triangular document-level pre-training on parallel data triplets**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7845–7858, Singapore. Association for Computational Linguistics.
- Hongyuan Lu, Haoyang Huang, Dongdong Zhang, Haoran Yang, Wai Lam, and Furu Wei. 2023. **Chain-of-Dictionary Prompting Elicits Translation in Large Language Models**. *arXiv e-prints*, page arXiv:2305.06575.
- Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, Alexandre Muzio, Saksham Singhal, Hany Hassan Awadalla, Xia Song, and Furu Wei. 2021. **DeltaLM: Encoder-Decoder Pre-training for Language Generation and Translation by Augmenting Pretrained Multilingual Encoders**. *arXiv e-prints*, page arXiv:2106.13736.
- Lesly Miculicich, Dhananjay Ram, Nikolaos Pappas, and James Henderson. 2018. **Document-level neural machine translation with hierarchical attention networks**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2947–2954, Brussels, Belgium. Association for Computational Linguistics.
- NLLB-Team. 2022. No language left behind: Scaling human-centered machine translation.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Machel Reid and Mikel Artetxe. 2022. **PARADISE: Exploiting parallel data for multilingual sequence-to-sequence pretraining**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 800–810, Seattle, United States. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. **Improving neural machine translation models with monolingual data**. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Zewei Sun, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Shujian Huang, Jiajun Chen, and Lei Li. 2022. **Re-thinking document-level neural machine translation**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3537–3548, Dublin, Ireland. Association for Computational Linguistics.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, page 3104–3112, Cambridge, MA, USA. MIT Press.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021a. **mT5: A massively multilingual pre-trained text-to-text transformer**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021b. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Jian Yang, Shuming Ma, Haoyang Huang, Dongdong Zhang, Li Dong, Shaohan Huang, Alexandre Muzio, Saksham Singhal, Hany Hassan, Xia Song, and Furu Wei. 2021a. [Multilingual machine translation systems from Microsoft for WMT21 shared task](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 446–455, Online. Association for Computational Linguistics.
- Jian Yang, Shuming Ma, Haoyang Huang, Dongdong Zhang, Li Dong, Shaohan Huang, Alexandre Muzio, Saksham Singhal, Hany Hassan, Xia Song, and Furu Wei. 2021b. [Multilingual machine translation systems from Microsoft for WMT21 shared task](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 446–455, Online. Association for Computational Linguistics.
- Jian Yang, Yuwei Yin, Shuming Ma, Haoyang Huang, Dongdong Zhang, Zhoujun Li, and Furu Wei. 2021c. [Multilingual agreement for multilingual neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 233–239, Online. Association for Computational Linguistics.
- Mingming Yang, Xing Wang, Min Zhang, and Tiejun Zhao. 2020. [Incorporating phrase-level agreement into neural machine translation](#). In *Natural Language Processing and Chinese Computing: 9th CCF International Conference, NLPCC 2020, Zhengzhou, China, October 14–18, 2020, Proceedings, Part I*, page 416–428, Berlin, Heidelberg. Springer-Verlag.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. [Hierarchical attention networks for document classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1480–1489, San Diego, California. Association for Computational Linguistics.
- Zhirui Zhang, Shuangzhi Wu, Shujie Liu, Mu Li, Ming Zhou, and Tong Xu. 2019. [Regularizing neural machine translation by target-bidirectional agreement](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI’19/IAAI’19/EAAI’19*. AAAI Press.
- Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2023. [Multilingual Machine Translation with Large Language Models: Empirical Results and Analysis](#). *arXiv e-prints*, page arXiv:2304.04675.

## A Number of Model Parameters

Model	Number of Parameters
M2M-100	418M
mBART	611M
MARGE	963M
mT5	1.23B*
DOCmT5	1.23B*
Baseline Model	862M
<b>BMA-SBT (Ours)</b>	862M

Table 5: Comparison in the number of parameters for the pre-trained models used in our experiments. \*: these models all use the model architecture of mT5-Large, and we report the number of model parameters taken from the original paper of mT5 reported by [Xue et al. \(2021b\)](#).

Table 5 presents the number of model parameters for the pre-trained models used in our experiments.